

The ideal Thurston-Andreev theorem and triangulation production

Gregory Leibon
Department of Mathematics
Dartmouth College
Hanover, NH

February 1, 2008

mailing address:

Gregory Leibon
29 Fletcher Circle
Hanover, NH
03755

e-mail address:

gleibon@dartmouth.edu

1 Introduction

The main result in this paper is a generalization of the convex ideal case of the Thurston-Andreev theorem when $\chi(M) < 0$. The proof naturally decomposes into its non-linear and linear aspects. The non-linear part is essentially a triangulation production theorem, stated in section 2.1. This theorem concerns taking a topological triangulation along with formal angle data and then “conformally flowing” this formal angle data to uniquely associated uniform angle data, where uniform angle data means the data contained in a geodesic triangulation of a hyperbolic surface. This flow is the gradient of an objective function related in a rather magical way to hyperbolic volume. That such a magical connection might exist was first explored in Bragger [1], and the hyperbolic volume needed in the case presented here was observed by my thesis advisor, Peter Doyle.

Conformally flowing turns out to be related to certain disk patterns and hyperbolic polyhedra, as discussed in section 2.2 and [3]. The linear part of this paper concerns gaining an explicit handle on which patterns can arise. In the end, a complete characterization of the “convex ideal” patterns in the $\chi(M) < 0$ case of the Thurston-Andreev theorem is presented, for the statement of which see section 2.3. It worth noting that in the torus and spherical cases that this generalization had already been accomplished. The toroidal case of the entire strategy used here has its origins in the beautiful and often overlooked work of Bragger [1] and can also be found in [4]. The spherical case of this generalization of the Andreev theorem was accomplished by Rivin in [5]. Whether there is proof directly using the spherical version of the techniques in this paper is still unknown. For an account of the theorem being generalized see Thurston’s [6].

This paper is organized as follows: section 2 contains the statements, set up, and notation needed to describe the theorems mentioned above. Section 3 contains the proof of the triangulation production theorem using a bit of hyperbolic geometry. Section 4 contains the proof of its corollary, the generalization of the Thurston-Andreev theorem stated in section 2.3; in the form of a “min flow max cut” type argument. In the final section a discussion of some known generalizations and some questions takes place.

I would like to thank my thesis advisor Peter Doyle for sharing his many beautiful ideas with me; without him the work here would not have been possible.

2 Statements and Notation




Figure 1: **A covector and a vector in geometric notation.** A vector will be denoted by placing its coefficients, $\{A^i\}$, in a copy of the triangulation with dashed lines and covector will contain its coefficients, $\{A_i\}$, in a copy of the triangular decomposition with solid lines. Thick lines will always denote a boundary edge, as in lower the left corner of pictured vector and covector. If in the picture we mean the non-specified values to be arbitrary we will surround the picture with a loop (as with the pictured vector) and if we mean the non-specified values to be zero the picture will not be surrounded (as with the pictured covector).

2.1 The Triangular Decomposition Theorem

Throughout this paper M will denote a compact two-dimensional surface with $\chi(M) < 0$. By geometry I will mean a hyperbolic structure. Uniqueness of geometries, triangulations and disk patterns is of course up to isometry.

The main theorem in this section really should be stated for the following structure, which generalizes the notion of triangulation.

Definition 1 *Let a triangular decomposition, \mathbf{T} , be a cell decomposition of M that lifts to a triangulation in M 's universal cover.*

We will keep track of the combinatorics of such a decomposition by denoting the vertices as $\{v_i\}_{i=1}^V$, the edges as $\{e_i\}_{i=1}^E$ and the triangles as $\{t_i\}_{i=1}^F$. Let E , V , F , ∂E and ∂V denote the sets of edges, vertices, faces, boundary edges, and boundary vertices respectively. For convenience this same notation will denote the cardinalities of these sets. Let $\{e \in S\}$ denote the set of edges on the surface in a collection of triangles S , and let $\{e \in v\}$ denote all the edges associated to a vertex v as if counted in the universal cover. The set of triangles containing a vertex $\{t \in v\}$ has the special name of the flower at v .

Note that in a triangular decomposition there are $3F$ slots $\{\alpha_i\}$ in which one can insert possible triangle angles, which we will place an order on and identify with a basis of a $3F$ dimensional real vector space. With this basis choice we will




Figure 2: **The geometric pairing of a vector and a covector** is achieved by placing the copy of the triangular decomposition corresponding to the vector on top of the triangular decomposition corresponding to the covector and multiplying the numbers living in the same angle slots.




Figure 3: For each vertex v the covector indicated in this figure will be denoted p^v

denote this vector space as \mathbf{R}^{3F} , and denote vectors in it as $x = \sum A^i \alpha_i$. Furthermore let α^i be a dual vector such that $\alpha^i(\alpha_j) = \delta_i^j$. With this we will view the angle at the slot α_i as $\alpha^i(x) = A^i$. It is rarely necessary to use this notation and instead to use the actual geometry as in figure 1. Notice the pairing of a vector and a covector denoted can be viewed geometrically as in figure 2. For a triangle t containing the angle slots α_i, α_j , and α_k let $d^t(x) = \{A^i, A^j, A^k\}$ and call $d^t(x)$ the angle data associated to t .

In order to live on an actual nonsingular geometric surface with geodesic boundary all such angles should be required to live in the subset of \mathbf{R}^{3F} where the angles at an interior vertex sum to 2π and the angles at a boundary vertex sum to π . Using the in figure 3, this encourages us to choose our possible angles in the affine flat

$$\mathbf{V} = \{x \in \mathbf{R}^{3F} \mid p^v(x) = 2\pi \text{ for all } v \in V - \partial V \text{ and } p^v(x) = \pi \text{ for all } v \in \partial V\}.$$

To further limit down the possible angle values we define the covector l^t as in




Figure 4: For each t we will denote the pictured covector as σ^t .

figure 4 Using the notation in figure 4, by the Gauss-Bonnet theorem we know that $k^t(x)$ defined as

$$k^t(x) = \sigma^t(x) - \pi$$

would be the curvature in a geodesic triangle with angle data $d^t(x)$. We will now isolate the open convex subset of \mathbf{V} where the curvature is negative and all angles are realistic.

Definition 2 *Let an angle system be a point in*

$$\mathbf{N} = \{x \in \mathbf{V} \mid k^t(x) < 0 \text{ for all } t \text{ and } \alpha^i(x) \in (0, \pi) \text{ for all } \alpha^i\}.$$

Note the actual angle data of a geodesic triangulation of a surface with negative curvature has its angle data living in this set.

Observe if $k^t(x) < 0$ then we may form an actual hyperbolic triangle with the angles in $d^t(x)$. For each $e \in t$ denote the length of the edge e with respect to this data as $l_t^e(x)$.

Suppose the triangles of \mathbf{T} fit together in the sense that $l_{t_1}^e(x) = l_{t_2}^e(x)$ whenever it makes sense. Then \mathbf{T} being a triangular decomposition implies every open flower is embedded in M 's universal cover and when the edge lengths all agree this flower can be given a hyperbolic structure which is consistent on flower overlaps. So we have formed a hyperbolic structure on M .

Definition 3 *Call an angle system u uniform if all the hyperbolic realizations of the triangles in u fit together to form a hyperbolic structure on M .*

In section 2.4 we will attempt to take a point in \mathbf{N} and deform it into a uniform point. Such deformations are located in an affine space and I will call them conformal deformations (see [3] for a more careful motivation of this terminology). To describe this affine space for each edge $e \in E - \partial E$ construct a vector w_e as in figure 5.




Figure 5: For each edge $e \in E - \partial E$ let w_e denote a vector as in this figure.

Definition 4 Using the notation from figure 5, a conformal deformation will be a vector in

$$C = \text{span}\{w_e \mid \text{for all } e \in E - \partial E\}.$$

Call x and y conformally equivalent if $x - y \in C$, and let

$$\mathbf{N}_x = (x + C) \cap \mathbf{N}$$

denote the conformal class of x .

The first thing worth noting is that if $x \in \mathbf{V}$ and y is conformally equivalent to x , then looking at the pairing between p^v with $v \in V - \partial V$ and the w_e (as in figure 3) we have

$$p^v(y) = p^v \left(x + \sum_{e \in \mathbf{P}} B^e w_e \right) = p^v(x) + \sum_{e \in \mathbf{P}} B^e p^v(w_e) = 2\pi + 0,$$

hence y is also in \mathbf{V} . Similarly for $v \in \partial V$.

To combinatorially understand the points in \mathbf{N} which we may conformally deform into uniform structures it is useful to express a particularly nasty set in the boundary of \mathbf{N} .

Definition 5 Let t be called a legal with respect to $x \in \partial \mathbf{N}$ if $d^t(x) = \{A^1, A^2, A^3\} \neq \{0, 0, \pi\}$ yet either $k^t(x) = 0$ or for some i we have $A^i = 0$. Let

$$\mathbf{B} = \{x \in \partial \mathbf{N} \mid x \text{ contains no legal } t\}.$$

We will prove the following theorem.

Theorem 1 If there is a uniform angle system conformally equivalent to x then it is unique, and for any angle system x with $(x + C) \cap \mathbf{B}$ empty there exists a conformally equivalent uniform angle system.

Much of what takes place here relies on certain basic invariants of conformal deformations, which is the subject of the next section.




Figure 6: For each triangle t and $e \in t$ we will let ψ_t^e denote the pictured covector.

2.2 Ideal Disk Patterns

In this section we introduce the disk patterns that a conformal class is related to. We begin with the introduction and interpretation of certain key conformal invariants. Using the notation in figure 6 for each edge $e \in \partial V$ we will denote ψ_t^e as ψ^e while for each edge $e \in E - \partial E$ associated with triangles t_1 and t_2 we will let

$$\psi^e = \psi_{t_1}^{e_1} + \psi_{t_2}^{e_1}.$$

We will call the ψ^e covector the formal angle complement at e . Let the formal intersection angle be defined as

$$\theta^e(x) = \pi - \psi^e(x)$$

when $e \in E - \partial E$. and

$$\theta^e(x) = \frac{\pi}{2} - \psi^e(x)$$

when $e \in \partial E$.

Looking at the pairing between a ψ^e and w_e (as in figure 2), we see if y is conformally equivalent to x then

$$\psi^e(y) = \psi^e \left(x + \sum_{f \in \mathbf{T}} B^f w_f \right) = \psi^e(x) + \sum_{f \in \mathbf{T}} B^f \psi^e(w_f) = \psi^e(x),$$

and indeed for each relevant edge e we see ψ^e and θ^e are conformal invariants.

The fundamental reason why theorem 1 is related to circle patterns and polyhedra construction is the following wonderful observation.

Observation 1 *At a uniform angle system u and $e \in E - \partial E$, $\theta^e(u)$ is the intersection angle of the circumscribing circles of the two hyperbolic triangles sharing e .*




Figure 7: In the figure we have our specified hyperbolic plane $H^2 \subset H^3$ realized as the intersection of the unit sphere at the origin with the xy -plane in \mathbf{R}^3 via the Poincaré disk model of H^3 . The point p in the figure is mapped to the point labeled x under the inversion I . We also are viewing a triangle on that plane and its associated prism in this model.

Proof: The proof given here relies on a principle fundamental to everything that takes place here, namely that questions taking place on a hyperbolic surface can often be best interpreted by viewing the question in three-dimensional hyperbolic space, H^3 . Let us call H^2 the hyperbolic plane in H^3 viewed as in figure 7. The inversion, I , through the sphere of radius $\sqrt{2}$ centered at the south pole interchanges our specified H^2 with the upper half of the sphere at infinity S_u^∞ . Notice when viewed geometrically this map sends a point $p \in H^2$ to the point where the geodesic perpendicular to H^2 containing p hits S_u^∞ (see figure 7). In particular being an inversion any circle in the xy -plane is mapped to a circle on the sphere at ∞ , S^∞ .

The use of this mapping will require the introduction of an object that will be crucial in proving theorem 1.

Definition 6 Place a hyperbolic triangle on a copy of $H^2 \subset H^3$. Let its associated prism be the convex hull of the set consisting of the triangle unioned with the geodesics perpendicular to this $H^2 \subset H^3$ going through the triangles vertices as visualized in figure 7. Denote the prism relative to the hyperbolic triangle constructed from the data $d = \{A, B, C\}$ as $P(d)$.

Now back to our proof. Let u be a uniform angle system and let t_1 and t_2 be a pair of triangles sharing the edge e . Place them next to each other in our H^2 from figure 7. Notice t_1 and t_2 have circumscribing circles in the xy -plane, which




Figure 8: The notation for an ideal prism associated to hyperbolic angle data $\{A, B, C\}$, viewed for convenience in the Klein model.

correspond to either circles, horocircles, or bananas in the H^2 geometry. Since the Poincare model is conformal the intersection angle of these circles is precisely the hyperbolic intersection angle. Being an inversion I is conformal, so these circles are sent to circles at infinity intersecting at the same angle and going through the ideal points of the neighboring $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$. But these circles at infinity are also the intersection of S^∞ with the spheres representing the hyperbolic planes forming the top faces of $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$. So the intersection angle of these spheres is precisely the sum of the angles inside $P(d_{t_1}(u))$ and $P(d_{t_2}(u))$ at the edge corresponding to e , which we will now see is $\theta^e(u)$. In fact we will show that this decomposition of the intersection angle is precisely the decomposition

$$\theta^e(u) = \theta_{t_1}^e(u) + \theta_{t_2}^e(u).$$

Now pick an i and let $d^{t_i}(u)$ be denoted by $\{A, B, C\}$. Assume our specified edge e corresponds to the a in figure 8. From figure 9 we see the angles in figure 8 satisfy the system of linear equation telling us that interior angles of the prism sum to π at each vertex of the prism. Solving this system for the needed angle, A^* , we find that indeed

$$A^* = \frac{\pi + A - B - C}{2} = \theta_{t_i}^e(u),$$

as claimed.

q.e.d




Figure 9: In this figure we have cut off an ideal vertex of a convex hyperbolic polyhedron with a horoball. Horospheres are flat planes in H^3 and our cut produces an Euclidean polygon in this flat plane. In particular the interior angles at an ideal vertex of a convex hyperbolic polyhedron are those of an Euclidean polygon. As indicated, this is best seen the upper half space model of H^3 , with the ideal vertex being sent to ∞ . (Nice applications of this observation can be found in [6].)

At this point it is useful to name the appropriate home of the possible intersection angle assignments. Just as we did with the angle slots, let the edges correspond to the basis vectors of an E dimensional vector space, which we will denote \mathbf{R}^E with this basis choice. We will be viewing this as the space of possible angle complements. Denote these vectors as $p = \sum_{e \in \mathbf{T}} E^e \psi_e$. let

$$\Psi : \mathbf{R}^{3F} \rightarrow \mathbf{R}^E$$

be the linear mapping given by

$$\Psi(x) = \sum_{e \in \mathbf{T}} \psi^e(x) \psi_e, \quad (1)$$

and note by observation 1 that we do indeed hit the intersection angle discrepancies when applying Ψ to an uniform angle system. Let ψ^e denote the algebraic dual to ψ_e . To justify this abuse of notation, note that $\psi^e(\Psi(x)) = \psi^e(x)$. Similarly let $\theta^e(p) = \pi - \psi^e(p)$ for $e \in E - \partial E$ and $\theta^e(p) = \frac{\pi}{2} - \psi^e(p)$ when $e \in \partial E$.

Definition 7 A point $p \in \mathbf{R}^E$ relative to a triangular decomposition \mathbf{T} will be called an ideal disk pattern if $p = \Psi(u)$ for some uniform angle system u .

In less dramatic terms this simply says that such a p can be realized geometrically. Theorem 1 provides us with the following corollary telling us such realizations are unique.




Figure 10: For each edge let m_e be the pictured vector.

Corollary 1 *For any $p \in \mathbf{R}^E$ at most one uniform u can satisfy $\Psi(u) = p$.*

Proof: Ψ has rank E since the pairing of ψ^e with the vector m_e in figure 10 satisfies $\Psi(m_e) = \psi_e$ for each edge e . By the conformal invariance noted in the previous section the null space contains the $E - \partial E$ dimension space C and is

$$3F - E = 2E - \partial E - E = E - \partial E$$

dimensional, so C is precisely the null space. In particular all angle systems which could conceivably hit a specified p will be in $\Psi^{-1}(p)$, which is $x + C$ for some x . So theorem 1 guarantees the uniqueness of the associated angle system.

q.e.d

In the next section we will derive affine conditions on a point $p \in \mathbf{R}^E$ relative to a triangular decomposition \mathbf{T} to be an ideal disk pattern.

I will finish this section by noting that such patterns are related to a certain family of hyperbolic polyhedra. Namely if we take the polyhedra constructed by placing a geodesic triangulation on an $H^2 \subset H^3$ and forming $\bigcup_{t \in \tilde{\mathbf{T}}} P(d_t(u))$. Note as a scholium to fact 1 that the dihedral angles in the polyhedra are precisely the intersection angles in the disk patterns. So any question about such ideal disk patterns can be translated into a question concerning such polyhedra (see [3] for details). In particular in the following sections we will be placing constraints on the θ^e to be in $(0, \pi]$, and its worth noting that this corresponds precisely to the convex case among these polyhedra. The results there can be viewed as a characterization of the possible dihedral angles in such convex polyhedra.

2.3 A Thurston-Andreev Type Theorem

In the previous section we learned that an ideal disk pattern is always unique when it exists, and we are now left to deal with the dilemma of finding good existence criteria. This will be accomplished by noting the restrictions placed on $\Psi(x)$ for $x \in \mathbf{N}$. For example the curvature assumptions in \mathbf{N} produces the following restrictions:

Sub-lemma 1 When $x \in \mathbf{N}$ we have $\psi_t^e(x) \in (-\frac{\pi}{2}, \frac{\pi}{2})$.

Proof: Let $d^t(x) = \{A, B, C\}$ and note since $B + C \leq A + B + C = \sigma^t(x) < \pi$ and $A < \pi$ we have

$$-\frac{\pi}{2} < -\frac{A}{2} < \psi_t^e = \frac{B + C - A}{2} < \frac{B + C}{2} < \frac{\pi}{2},$$

as needed.

q.e.d

In this section we will be strengthening this restriction to the strict convex case when the ideal disk pattern satisfies $p \in (0, \pi)^{(E-\partial E)} \times (0, \frac{\pi}{2})^{\partial E}$. An angle system producing such an ideal disk pattern will live in

$$\mathbf{D} = \left\{ x \in \mathbf{N} \mid \Psi(x) \in (0, \pi)^{(E-\partial E)} \times \left(0, \frac{\pi}{2}\right)^{\partial E} \right\},$$

and let us call set this the set of Delaunay angle systems. Aside from arising naturally in practice from the Delaunay triangulation of a set of points (see [3]), these angle systems are remarkably easy to work with because they completely eliminate the possibility of a conformally equivalent badly behaved sets of angles.

Lemma 1 $x \in \mathbf{D}$ is never conformally equivalent to a point in $\partial \mathbf{N}$ where for some triangle $d^t(x) = \{0, 0, \pi\}$.

Proof: To see this fact assume to the contrary that for some t and c we have $d^t(x + c) = \{0, 0, \pi\}$. Let e be the edge of t across from t 's π and let t_1 be t 's neighbor next to e if it exists.

Note by sublemma 1 that the conformally invariant $\psi^e(x) \in (0, \pi)$ would (even in the best possible case when e is not on the boundary) have to satisfy the inequality $\psi^e(x + C) = -\frac{\pi}{2} + \psi_{t_1}^e \leq 0$, contradicting the fact $x \in \mathbf{D}$.

q.e.d

The elimination of such possibilities allows us to immediately apply theorem 1, and arrive at...

Corollary 2 Every point of \mathbf{D} has a unique ideal disk pattern associated to it.

Proof: Let $x \in \mathbf{D}$ and note we are searching for a uniform angle system in $x + C \cap \mathbf{N} = \Psi^{-1}(\Psi(x)) \cap \mathbf{N}$. From lemma 1 we have that if $x \in \mathbf{D}$ then no element in $x + C$ could possibly be in B and the corollary follows from theorem 1.

q.e.d

The goal at this point is to determine necessary and sufficient conditions on a point in \mathbf{R}^E to insure that it is $\Psi(x)$ for some Delaunay angle system. We immediately have that any such point is $p \in (0, \pi)^{(E-\partial E)} \times (0, \frac{\pi}{2})^{\partial E}$, and our first non-trivial necessary condition is the condition related to the fact that the angles at the internal vertex in a geometric triangulation sum to 2π and at a boundary vertex sum to π .

$$(n_1) \begin{cases} \sum_{e \in v} \psi^e(p) = 2\pi & \text{if } v \in V - \partial V \\ \sum_{e \in v} \psi^e(p) = \pi & \text{if } v \in \partial V \end{cases}$$

Proof: To see the necessity of (n_1) we will show

$$\Psi(V) = \{p \in \mathbf{R}^E \mid p \text{ satisfies } (n_1)\}.$$

First note that if $p = \Psi(x)$ then

$$\sum_{e \in v} \psi^e(p) = \sum_{e \in v} \psi^e(x) = p^v(x).$$

So by choosing $x \in V$ we see $\Psi(V)$ is included in

$$W = \{p \in \mathbf{R}^E \mid p \text{ satisfies } (n_1)\}.$$

Recall from the proof of corollary 1 that $\Psi(\mathbf{R}^{3F}) = \mathbf{R}^E$. So we may express any $p \in W$ as $p = \Psi(x)$ and the above computation guarantees $x \in V$ as needed.

q.e.d

The second necessary condition is a global one; namely an insistence that for every set S of $|S|$ triangles in \mathbf{T} that

$$(n_2) \quad \sum_{e \in S} \theta^e(p) > \pi |S|.$$

Proof: Verifying (n_2) 's necessity relies on the following formula.

Formula 1 *Given a set of triangles S*

$$\sum_{\{e \in S\}} \theta^e(x) = \sum_{t \in S} \left(\pi - \frac{k^t(x)}{2} \right) + \sum_{e \in \partial S - \partial E} \left(\frac{\pi}{2} - \psi_t^e(x) \right),$$

with the t in $\psi_t^e(x)$ term being the triangle on the non- S side of e .

$$\begin{aligned}
\sum_{\{e \in S\}} \theta^e(x) &= \sum_{e \in S - \partial E} (\pi - \psi^e(x)) + \sum_{e \in S \cap \partial E} \left(\frac{\pi}{2} - \psi^e(x) \right) \\
&= \sum_{e \in S - \partial E} \left(\left(\frac{\pi}{2} - \psi_{t_1}^e(x) \right) + \left(\frac{\pi}{2} - \psi_{t_2}^e(x) \right) \right) + \sum_{e \in S \cap \partial E} \left(\frac{\pi}{2} - \psi^e(x) \right) \\
&= \sum_{t \in S} \left(\pi + \frac{\pi - \sigma^t(x)}{2} \right) + \sum_{e \in \partial S - \partial E} \left(\frac{\pi}{2} - \psi_t^e(x) \right)
\end{aligned}$$

with the t in $\psi_t^e(x)$ term being the triangle on the non- S side of e . Substituting the definition of $k^t(x)$ gives the needed formula.

To apply this formula note for any point $x \in \mathbf{N}$ that $-k^t(x) > 0$ and from sublemma 1 that $\frac{\pi}{2} - \psi_t^e(x) > 0$. So removing these terms from the above formula strictly reduces its size and when summed up we arrive at (n_2) .

q.e.d

With these two necessary conditions we arrive at a pattern existence theorem (see section 5 for a stronger result):

Theorem 2 *If*

$$p \in D = \left\{ q \in (0, \pi)^{(E - \partial E)} \times \left(0, \frac{\pi}{2} \right)^{\partial E} \mid q \text{ satisfies } (n_1) \text{ and } (n_2) \right\}$$

then p is an ideal disk pattern.

By corollary 2 above and the fact that (n_1) and (n_2) are necessary this theorem would follow if we knew

$$\Psi(\mathbf{D}) = D.$$

Notice the fact that (n_1) and (n_2) are necessary guarantees that $\Psi(\mathbf{D}) \subset D$, and we left to explore Ψ 's surjectivity. It is this bit of linear algebra and will be dealt with in section 4.

3 The Non-linear Argument

The non-linear argument is dependent on a rather remarkable link between hyperbolic volume and uniform structures. Let the volume of the prism $P(d^t(y))$ be denoted $V(d^t(y))$, we will be exploring the objective function

$$H(y) = \sum_{t \in \mathbf{T}} V(d^t(y))$$

on \mathbf{N}_x . First let us note this function has the correct objective.

Fact 1 y is a critical point of H on \mathbf{N}_x if and only if y is uniform.

Proof: To see this fact requires an understanding of H 's differential at each x . As usual for a function on a linear space like \mathbf{R}^{3F} we use translation to identify the tangent and cotangent spaces at every point with \mathbf{R}^{3F} and $(\mathbf{R}^{3F})^*$, and express our differentials in the chosen basis.

Lemma 2 $dH(z) = \sum_{t \in \mathbf{T}} (\sum_{e \in t} h_t^e(z) d\theta_t^e)$ with the property that $h_t^e(z)$ uniquely determines the length $l_t^e(z)$.

This formula will be proved in section 2, for now lets see how to use it. Recall that $T_z(\mathbf{N}_x) = \text{span}\{w_e \mid e \in E - \partial E\}$ from definition 4. Now simply observe that $\theta_t^e(w_f) = \pm \delta_f^e$ with the sign depending on whether t contains the negative of positive half of w_e . So at a critical point y we have

$$0 = dH(y)(w_e) = h_{t_1}^e(y) - h_{t_2}^e(y)$$

where t_1 and t_2 are the faces sharing e . So from the above lemma and the fact that the w_e span $T_y(\mathbf{N}_x)$ we have that $l_{t_1}^e(y) = l_{t_2}^e(y)$ for each edge is equivalent to y being critical, as needed.

q.e.d

We may now prove the uniqueness assertion in theorem 1.

Proposition 1 If \mathbf{N}_x contains a uniform angle system this angle system is unique in \mathbf{N}_x .

Proof: The following lemma will be proved in section 3.2.

Lemma 3 H is a strictly concave, smooth function on \mathbf{N}_x and continuous on $\overline{\mathbf{N}}_x$.

Now a smooth strictly concave function like E has a most one critical point in any open convex set, which proves proposition 1.

q.e.d

Now its time to explore the existence of critical points. Given a pre-compact open set U and a continuous function F on \overline{U} we automatically achieve a maximum. For this maximum to be a critical point it is enough to know that F is differentiable in U and that the point of maximal F is in U .

One way to achieve this is to show that for any boundary point y_0 that there is a direction v , an $\epsilon > 0$ and a $c > 0$ such that $l(s) = y_0 + sv$ satisfies

$$l(0, \epsilon) \subset U$$

and

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} F(l(s)) > c,$$

for all $s \in (0, \epsilon)$. This follows since under these hypotheses $F(l(s))$ is continuous and increasing on $[0, \epsilon)$ and y_0 certainly could not have been a point where F achieved its maximum.

In our setting we have the following lemma to be proved in section 3.3.

Lemma 4 *For every point y_0 in $\partial\mathbf{N}$ and every direction v such that*

$$l(0, \infty) \cap \mathbf{N} \neq \emptyset$$

and

$$l[0, \infty) \cap \mathbf{B} = \emptyset$$

we have

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} H(l(s)) = \infty.$$

By convexity of \mathbf{N}_x for each boundary point there is such a v , so by the above observations we now have that H achieves its unique critical in \mathbf{N}_x , as needed to prove theorem 1.

3.1 The Differential: The Computation of Lemma 2

In this section we gain our needed understanding of the differential as expressed in lemma 2. To get started note the sum in dH is over all triangles but the fact concerns only each individual one. So we may restrict our attention to one triangle. One way to prove lemma 2 is to explicitly compute a formula for the volume in terms of the Lobachevsky function and then find its differential. This method can be found carried out in [3]. Here we present an argument using Schlaflis formula for volume deformation. This technique has a wider range of application as well as being considerably more interesting.

To start with we will recall Schlafli's formula for a differentiable family of compact convex polyhedra with fixed combinatorics. Let $\{edges\}$ denote the set of edges and $l(e)$ and $\theta(e)$ be the length and dihedral angle functions associated to an edge e . Schlafli's formula is the following formula for the deformation of the volume with in this family

$$dV = -\frac{1}{2} \sum_{edges} l(e) d(\theta(e)).$$

In the finite volume case when there are ideal vertices the formula changes from measuring the length of edges $l(e)$ to measuring the length of the cut off edges $l^*(e)$. Let us now recall how $l^*(e)$ is computed. First fix a horosphere at each ideal vertex. Then note from any horosphere to a point and between any pair of horospheres there is a unique (potentially degenerate) geodesic segment perpendicular to the horosphere(s). $l^*(e)$ is the signed length of this geodesic segment; given a positive sign if the geodesic is out side the horosphere(s) and a negative sign if not. Schläfli's formula is independent of the horosphere choices in this construction, and I will refer to this fact as the horoball independence principle. It is worth recalling the reasoning behind this principle, since the ideas involved will come into play at several points in what follows.

The Horosphere Independence Reasoning: Recall from the proof of observation 1 that at an ideal vertex v we have the sum of the dihedral angles satisfying $\sum_{e \in v} \theta(e) = (n - 2)\pi$, and in particular

$$\sum_{e \in v} d\theta(e) = 0.$$

Looking at figure 9 we see by changing the horosphere at the ideal vertex v that $l^*(e)$ becomes $l^*(e) + c$ for each $e \in v$ with c a fixed constant. Hence by our observation about the angle differentials

$$-2dV = \sum_{edges} l^*(e)d\theta_e = \sum_{\{e \in v\}^c} l^*(e)d\theta(e) + \sum_{e \in v} (l^*(e) + c)d\theta(e)$$

and dV is seen to be independent of the horosphere choices.

q.e.d

Now let us look at our prism. Let the notation for the cut off edge lengths coincide with the edge names in figure 8. Since we may choose any horospheres let us choose those tangent to the hyperbolic plane which our prism is symmetric across. In this case note the lengths of $(ab)^*$, $(bc)^*$ and $(ac)^*$ are zero. Recalling from the proof of observation 1 that

$$A^* = \frac{\pi + A - B - C}{2}$$

and viewing $V(d^t(x))$ as a function on

$$\{(A, B, C) \in (0, \pi)^3 : 0 < A + B + C < \pi\}$$

we see from Schläfli's formula that

$$dV = -a^*dA^* - b^*dB^* - c^*dC^*.$$

Note that lemma 2 will follow from the following formula.




Figure 11: The face of our prism containing a along with the horocircle slices of the horospheres tangent to the hyperbolic plane through which our prism is symmetric.

Formula 2

$$a^* = 2 \ln \left(\sinh \left(\frac{a}{2} \right) \right).$$

Proof: To begin this computation look at the face of the prism containing a as in figure 11. Notice this face is decomposed into four quadrilaterals as in figure 12. Note that just as with the above reasoning concerning the independence of horosphere choice we have an independence of horocircle choice and

$$\frac{a^*}{2} = (t^* - h^*) - (h^* - s^*).$$

In fact $t^* - h^*$ and $h^* - s^*$ are independent of this horocircle choice as well and it is these quantities we shall compute.

Look at the figure 12 and notice using the horocircle tangent to the $\frac{a}{2}$ geodesic that $h^* - s^*$ becomes precisely h^* . Viewing this situation as in figure 13 we can now read off from figure 13 that

$$h^* - s^* = -\ln \left(\operatorname{sech} \left(\frac{a}{2} \right) \right).$$

Similarly notice that

$$-h^* + t^* = \ln(\operatorname{sech}(l)),$$




Figure 12: Pictured here is one of the four triply right angled quadrilaterals from figure 11. Such quadrilaterals are known as Lambert quadrilaterals and it is a well know relationship that $\tanh\left(\frac{a}{2}\right) = \operatorname{sech}(l)$, see for example [2]. (In fact it follows nearly immediately from the perhaps better known Bolya-Lobachevsky formula.)

which as observed in figure 12 implies

$$-h^* + t^* = \ln\left(\tanh\left(\frac{a}{2}\right)\right).$$

With these computations we now have

$$a^* = 2\left(\ln\left(\tanh\left(\frac{a}{2}\right)\right) - \ln\left(\operatorname{sech}\left(\frac{a}{2}\right)\right)\right) = 2\ln\left(\sinh\left(\frac{a}{2}\right)\right)$$

as needed.

q.e.d

3.2 Convexity: The proof of Lemma 3

To prove H is strictly concave we start with the observation that the objective function H will certainly be a strictly concave function on \mathbf{N}_x if the prism volume function $V(d^t(x))$ viewed as a function on

$$\{(A, B, C) \in (0, \pi)^3 : A + B + C < \Pi\}$$

turned out to be strictly concave. In fact it is worth noting that this implies H is then strictly concave on all of $(0, \pi)^{3F}$ (see section 5.3).




Figure 13: Here we have placed the lower triangle from figure 12 into the upper-half plane model sending the ideal vertex to infinity and the $\frac{a}{2}$ segment onto the unit circle as pictured. Recall that the unit circle in this picture can be parameterized by hyperbolic distance from i via $\tanh(s) + \operatorname{sech}(s)i$.

There are several nice methods to explore the concavity of $V(A, B, C)$. One could simply check directly that V 's Hessian is negative definite (as done in [3]), or one could exploit the visible injectivity of the gradient, or one could bootstrap from the concavity of the ideal tetrahedron's volume. It is this last method that will be presented here. The crucial observation is that any family of ideal prism can be decomposed into three ideal tetrahedra as in figure 14. So we have

$$V(A, B, C) = \sum_{i=1}^3 T_i(A, B, C),$$

where T_i is the volume of the i^{th} tetrahedra in this decomposition.

Let us note some properties of the ideal tetrahedra and its volume. First recall from figure 9 that the dihedral angles corresponding to the edges meeting at a vertex of an ideal tetrahedron are the angles of a Euclidean triangle. In particular the constraints

$$\sum_{e \in v} \theta^e = \pi$$

at each vertex guarantee that an ideal tetrahedron is uniquely determined by any pair of dihedral angles α and β corresponding to a pair of edges sharing a vertex.




Figure 14: A decomposition of the ideal prism into three ideal tetrahedra. Notice that the angles in this decomposition are determined by the affine conditions coming from the ideal vertices (see figure 9) along with the condition that the angles meeting along an edge slicing a prism face sum to π . In particular all the angles depend affinely on the angles $\{A, B, C\}$.

Further more any pair of angles in

$$\{(\alpha, \beta) : \alpha + \beta < \pi\}$$

determines an ideal tetrahedron. Note the following fact (see [4]).

Fact 2 *The ideal tetrahedrons volume function, $T(\alpha, \beta)$, is strictly concave on the set*

$$\{(\alpha, \beta) : \alpha + \beta < \pi\}$$

and continuous on this set's closure.

From figure 14 each of the α_i and β_i of the i^{th} tetrahedron depend on the (A, B, C) affinely. So this fact immediately provides us with the continuity assertion in lemma 3. To exploit the tetrahedron's concavity we will use the following lemma.

Lemma 5 *Let T be a strictly concave function on the convex set $U \subset \mathbf{R}^m$ and for each i let L_i be an affine mapping from \mathbf{R}^n to \mathbf{R}^m taking the convex set V into U . Then the function*

$$V(\vec{x}) = \sum_{i=1}^k T(L_i(\vec{x}))$$

is strictly concave on V provided $L_1 \times \dots \times L_k$ is injective.

Proof: Let $l(t)$ be a line such that $l(0) = a \in V$ and $l(1) = b \in V$; and let $t \in (0, 1)$. Note by the concavity of T that

$$V(l(t)) = \sum_{i=1}^3 T(L_i(l(t))) \geq$$

$$\sum_{i=1}^k (T(L_i(l(a))) + t(T(L_i(l(b))) - T(L_i(l(a)))) = V(l(a)) + t(V(l(b)) - V(l(a))).$$

Since T is in fact strictly concave the inequality would be strict if for some i we knew $L_i(l(t))$ was a non-trivial line. Fortunately this is guaranteed by the injectivity of $L_1 \times \dots \times L_k$, and we are done.

q.e.d

Letting

$$L_i(A, B, C) = (\alpha_i(A, B, C), \beta_i(A, B, C))$$

we see that V we will satisfy the lemma if, for example, the mapping

$$(\alpha_1(A, B, C), \alpha_2(A, B, C), \alpha_3(A, B, C))$$

is injective. Looking at the decomposition in figure 14 we see that we may in fact choose $\alpha_1(A, B, C) = A$, $\alpha_2(A, B, C) = B$, and $\alpha_3(A, B, C) = C$. So indeed, we have our required injectivity and $V(A, B, C)$ is strictly concave as needed.

3.3 Boundary Control: Proof of Lemma 4

Before proving lemma 4 we will rephrase it slightly. Namely note that the compactness of \overline{U} guarantees that $l(s)$ eventually hits the boundary again at some y_1 for some unique $s > 0$. So we may change the speed of our line and assume we are using the line connecting the two boundary points, namely

$$l(s) = (1 - s)y_0 + sy_1.$$

So lemma 4 is equivalent to the following lemma.

Lemma 6 *For every pair of points y_0 and y_1 in $\partial\mathbf{N}$ but not in B with $l(s) \cap \mathbf{N} \neq \emptyset$ we have*

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} H(l(s)) = \infty.$$

Proof: Recalling that $H(d^t(x)) = \sum_{t \in \mathbf{T}} V(d^t(x))$ we see the lemma will follow if we demonstrate that for any triangle

$$-\infty < \lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) \leq \infty,$$

and for some triangle

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) = \infty.$$

The boundary is expressed in terms of angle data, so it would be nice to express the $-2 \ln(\sinh(\frac{a}{2}))$ coefficient in front of the dA^* in dV (as computed in section 3.1) in terms of the angle data. In fact we can do even better and put this term in a form conveniently decoupling the angle and curvature.

Formula 3 $-2 \ln(\sinh(\frac{a}{2}))$ is equal to

$$\ln(\sin(B)) + \ln(\sin(C)) - \ln\left(\frac{\cos(A - k^t(x)) - \cos(A)}{k^t(x)}\right) - \ln(-k^t(x)).$$

This formula relies only on the hyperbolic law of cosines which tells us

$$\cosh(a) = \frac{\cos(B) \cos(C) - \cos(A)}{\sin(A) \sin(B)}.$$

Using this relationship and the definition of $k^t(x)$ we now have

$$-2 \ln\left(\sinh\left(\frac{a}{2}\right)\right) = -\ln\left(\sinh^2\left(\frac{a}{2}\right)\right) = -\ln\left(\frac{\cosh(a) - 1}{2}\right)$$

$$= -\ln \left(\frac{\cos(B+C) + \cos(A)}{\sin(B)\sin(C)} \right) = -\ln \left(\frac{-\cos(A - k^t(x)) + \cos(A)}{\sin(B)\sin(C)} \right),$$

as needed.

Using this formula we will now enumerate the possible y_0 and the behavior of $\frac{d}{ds} V(d^t((s)))$ in these various cases. Let C denote a finite constant. We will be using the fact that if $L(s)$ is an affine function of s satisfying $\lim_{s \rightarrow 0^+} L(s) = 0$ then $\lim_{s \rightarrow 0^+} \ln |\sin(L(s))|$ and the $\lim_{s \rightarrow 0^+} \ln |D(L(s))|$ can both be expressed as $\lim_{s \rightarrow 0^+} \ln(s) + C$. Furthermore for convenience let $d^t(y_i) = \{A_i, B_i, C_i\}$.

1. When $d^t(y_0)$ contains no zeros and $k^t(y_0) \neq 0$ we have that $\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s)))$ is finite.
2. When $d^t(y_0) = \{0, 0, \pi\}$, we have

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) = \frac{1}{2} \lim_{s \rightarrow 0^+} \ln(s)(\sigma^t(y_1 - y_0)) - \frac{1}{2} \lim_{s \rightarrow 0^+} \ln(s)(\sigma^t(y_1 - y_0)) + C = C.$$

3. In the case where $d^t(y_0)$ contains zeros but $k^t(y_0) \neq 0$ for each zero (assumed to be A_0 below) we produce a term in the form

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) = \lim_{s \rightarrow 0^+} (-\ln(s)(A_1 - A_0)),$$

plus some finite quantity.

4. When $k^t(y_0) = 0$ and no angle is zero

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) = \lim_{s \rightarrow 0^+} \frac{1}{2} \ln(s)(\sigma^t(y_1 - y_0)) + C.$$

5. When $k^t(y_0) = 0$ and one angle, say A_0 , in $d^t(y_0)$ is zero we have

$$\lim_{s \rightarrow 0^+} \frac{d}{ds} V(d^t((s))) = -2 \lim_{s \rightarrow 0^+} \ln(s)(A_1 - A_0) + \lim_{s \rightarrow 0^+} \ln(s)(\sigma^t(y_1 - y_0)) + C.$$

So the first two cases produce finite limits. In order to understand the next three limits we make some simple observations. First if $A_0 = 0$ and $l(s) \cap \mathbf{N} \neq \emptyset$ then $A_1 - A_0 > 0$. So limits from the third case evaluate to $+\infty$. Secondly note that when $k^t(y_0) = 0$ and $l(s) \cap \mathbf{N} \neq \emptyset$ that $\sigma^t(y_1 - y_0) = A_1 + B_1 + C_1 - (A_0 + B_0 + C_0) < 0$ and hence the limits from the fourth case are $+\infty$ as well. Combining these observations we see the fifth case always produces a $+\infty$ limit as well.

So for each triangle the answer is indeed finite or positive infinity. So all we need to do is guarantee that for some triangle we achieve $+\infty$. To do this note that

in order for y_0 to be on the boundary of \mathbf{N} and not in B there is some triangle t such that $d^t(y_0) = \{A_0, B_0, C_0\} \neq \{0, 0, \pi\}$ however, either $k^t(y_0) = 0$ or some angle is zero. So we have at least one triangle in case 3,4, or 5 as needed.

q.e.d

It is worth noting that the choice of the terminology bad for the set \mathbf{B} is due to the fact that at such a point all triangles would fall into cases one or two above, and in the process we lose our needed control over H .

4 The Linear Argument

To see the surjectivity of Ψ from \mathbf{D} to D let us assume the contrary that $\Psi(\mathbf{D})$ is strictly contained in D and produce a contradiction. With this assumption we have a point p on the boundary of $\Psi(\mathbf{D})$ inside D . Note $p = \Psi(y)$ for some $y \in \partial\mathbf{D}$. Furthermore note $(C + y) \cap \mathbf{D}$ is empty, since otherwise for some $w \in C$ we would have $(y + w) \in \mathbf{D}$ which along with the fact that Ψ is an open mapping when restricted to V would force $p = \Psi(y) = \Psi(y + w)$ to be in the interior of $\Psi(\mathbf{D})$.

At this point we need to choose a particularly nice conformally equivalent version of y , which requires the notion of a stable boundary point of \mathbf{D} . Before defining stability, note since \mathbf{D} is a convex set with hyperplane boundary if $x \in \partial\mathbf{D}$ such that $(x + C) \cap \mathbf{D} = \emptyset$, then $(x + C) \cap \partial\mathbf{D}$ is its self a convex k dimensional set for some k .

Definition 8 *A point in $x \in \partial\mathbf{D}$ is stable if $(x + C) \cap \mathbf{D} = \emptyset$ and x is in the interior of $(x + C) \cap \partial\mathbf{D}$ as a k dimensional set. Any inequality forming \mathbf{D} violated in order to make x a boundary point will be called a violation.*

The key property of a stable point is that a conformal change $w \in C$ has $x + \epsilon w \in \overline{\mathbf{D}}^C$ for all $\epsilon > 0$ or for any sufficiently small $\epsilon > 0$ we have $x + \epsilon w$ must still be on $\partial\mathbf{D}$ and experience exactly the same violations as x . The impossibility of any other phenomena when conformally changing a stable point is at the heart of the arguments in lemma 7 and lemma 8 below. At this point surjectivity would follow if for a stable $x \in \partial\mathbf{D}$ we knew that $\Psi(x)$ could not be in D , producing the needed contradiction to our $p = \Psi(x)$ choice.

We will prove this by splitting up the possibilities into the two cases in lemma 7 and lemma 8.

Lemma 7 *If $x \in \partial\mathbf{D}$ is stable and $\alpha^i(x) = 0$ for α^i in some triangle where $k^t(x) < 0$, then $\Psi(x)$ is not in D .*




Figure 15: Here we have an accordion, the set of edges forming a loop indicated with the squiggly line. Notice to geometrically realize the zero angles in the pictures corresponds to squeezing the accordion. The algebraic inability to squeeze various accordions is at the heart of the Ψ 's surjectivity.

Proof: Look at an angle slot which is zero in triangle t_0 satisfying $k^{t_0}(x) < 0$. View this angle as living between the edges e_0 and e_1 . Note that in order for x to be stable that either e_1 is a boundary edge or the ϵw_{e_1} transformation (with its positive side in t_0) must be protected by a zero on the $-\epsilon$ side forcing the condition that $x + \epsilon w_{e_1} \in \bar{\mathbf{D}}^c$, or else for small enough ϵ we would have $x + \epsilon w_{e_1}$ being a conformally equivalent point on $\partial \mathbf{D}$ with fewer violations. When e_1 is not a boundary edge call this neighboring triangle t_1 and when it is a boundary edge stop this process. If we have not stopped let e_2 be another edge bounding a zero angle slot in t_1 and stop if it is a boundary edge. If it is not a boundary edge then there are two possibilities. If $k^{t_1} < 0$ repeat the above procedure letting e_1 play the role of e_0 and e_2 the role of e_1 and constructing an e_3 in a triangle t_2 . If $k^{t_1}(x) = 0$ conformally change x to

$$x + \epsilon w_{e_1} + \epsilon w_{e_2}.$$

Notice no triangle with $k^t(x) = 0$ can have two zeros by lemma 1, so for the initial zero violation to exist there must be a zero on the $-\epsilon$ side of ϵw_{e_2} . Once again we have determined an e_3 and t_2 .

Using this procedure to make our decisions we may continue this process forming a set of edges $\{e_i\}$ with the angle between e_i and e_{i+1} , $A^{i,i+1}(x)$, always equal to zero. I'll call such a set an accordion, see in figure 15 for an example. Since there are a finite number of edges an accordion either stops at a boundary edge or the accordion is endless. If the accordion is endless then eventually the sequence $\{e_i\}_1^\infty$ will have some $k < l$ such that $e_k = e_l$ and $e_{k+1} = e_{l+1}$, see figure 15.

(This is true by the pigeon hole principle since some edge e will appear an infinite number of times in this list and among its infinite neighbors there must be a repeat). We can produce a contradiction to this occurring. To do it first note if e_i and e_{i+1} are in t_i then $A^{i,i+1}(x) = \psi_{t_i}^{e_i} + \psi_{t_i}^{e_{i+1}}$. So for the set of edges $\{e_i\}_{i=k}^{l-1}$ we have

$$0 = \sum_{i=k}^{l-1} A^{i,i+1} = \sum_{i=k}^{l-1} \psi^{e_i}(x) > 0$$

our needed contradiction.

In the case the sequence did hit the boundary perform the accordion construction in the opposite direction. If we don't stop in this direction we arrive at the same contradiction. If we did then this computation still produces a contradiction on the accordion with the two boundary edges, since for a boundary edge in the triangle t we have $\psi_t^e(x) = \psi^e(x) \in (0, \frac{\pi}{2})$.

q.e.d

Lemma 8 *If a stable x satisfies the condition that if $\alpha^i(x) = 0$ then α_i is in a triangle t with $k^t(x) = 0$, then $\Psi(x)$ is not in D .*

Proof: In this case, in order for x to be a boundary point of \mathbf{D} for some t we have that $k^t = 0$. We will be looking at the nonempty set of all triangles with $k^t = 0$, Z . The first observation needed about Z is that it is not all of M and has a non-empty internal boundary (meaning $\partial Z - \partial M$). To see this note

$$\begin{aligned} \sum_{t \in \mathbf{T}} k^t(x) &= \sum_{e \in v} A^i - \pi F = \pi \partial V + 2\pi(V - \partial V) - \pi F \\ &= 2\pi V - (\pi \partial V + 3\pi F) + 2\pi F = 2\pi V - 2\pi E + 2\pi F = 2\pi \chi(M) < 0, \end{aligned}$$

so there is negative curvature somewhere.

By the stability of x once again there can be no conformal transformation capable of moving negative curvature into this set. Suppose we are at an internal boundary e_0 edge of Z , call the triangle on the Z side of the boundary edge t_0 and the triangle on the non-boundary edge t_{-1} . Since t_{-1} has negative curvature the obstruction to the ϵw_{e_0} transformation being able to move curvature out of Z must be due to t_0 . In order for t_0 to protect against this there must be zero along e_0 on the t_0 side.

Now we will continue the attempt to suck curvature out with a curvature vacuum. Such a vacuum is an element of C indexed by a set of Z edges. The key observation in forming this vacuum is once again lemma 1 telling us if an angle in t is zero and $k^t(x) = 0$ then there is only one zero angle in t . Let e_1 be the other




Figure 16: In this figure the region of zero curvature, Z , is the set with the bold boundary as labeled. We have also attached to Z all its vacuums and indicated them with the squiggly lines. As we can see in the figure these vacuum transformations fail to suck curvature out of Z , just as in the proof of lemma 8. Notice the set S , Z minus its vacuums, is in this case the pair of shaded triangles.

edge sharing the unique zero angle along e_0 in t_0 and if e_1 is another boundary edge we stop. If e_1 is not a boundary edge use $\epsilon(w_{e_1} + w_{e_0})$ to continue the effort to remove curvature. Continuing this process forms a completely determined set of edges and triangles, $\{e_i\}$ and $\{t_i\}$, and a sequence of conformal transformations $\epsilon \sum_{i=0}^n w_{e_i} \in C$, see figure 16.

We will now get some control over this vacuum. Note a vacuum never hits itself since if there is a first pair $k < l$ such that $t_k = t_l$ then t_k would have to have two zeros and zero curvature, which lemma 1 assures us is impossible. So any vacuum hits a boundary edge or pokes through Z into Z^c .

In fact with this argument we can arrive at the considerably stronger fact that two vacuums can never even share an edge. To see this, call a vacuum's side boundary any edge of a triangle in the vacuum facing a zero. Now simply note if the intersection of two vacuums contains an edge then it contains a first edge e_i with respect to one of the vacuums. There are two possibilities for this edge. One is that t_{i+1} has two zeros and $k^t(x) = 0$, which we showed was impossible in the previous paragraph. The other is that e_i is a side boundary of both vacuums. In this case we have an edge facing zero angles in both directions in triangles with zero curvature, so this would force $\psi^e(x) = \pi$, a contradiction. So either case is impossible, and indeed no distinct vacuums share an edge.

Let S be the removal from Z of all these vacuums, see figure 16. First I'd like to note that S is non-empty. Note every vacuum has side boundary. Since vacu-

ums cannot intersect themselves or share edges with distinct vacuums, S would be nonempty if side boundary had to be in Z 's interior. Look at any side boundary edge e of a fixed vacuum. Note e cannot be on $\partial Z - \partial M$ since then the vacuum triangle it belonged to would have at least two zeros and $k^t(x) = 0$. Furthermore, e cannot be on ∂M since then $\psi^e(x) = \frac{\pi}{2}$. So indeed S is nonempty.

Now let's observe the following formula.

Formula 4 *Given a set of triangles S*

$$\sum_{\{e \in S\}} \theta^e(x) = \sum_{t \in S} \left(\pi - \frac{k^t(x)}{2} \right) + \sum_{e \in \partial S - \partial M} \left(\frac{\pi}{2} - \psi_t^e(x) \right),$$

with the t in the $\psi_t^e(x)$ term being the triangle on the non- S side of e .

Proof:

$$\begin{aligned} \sum_{\{e \in S\}} \theta^e(x) &= \sum_{e \in S} (\pi - \psi^e(x)) \\ &= \sum_{e \in S - \partial M} \left(\left(\frac{\pi}{2} - \psi_{t_1}^e(x) \right) + \left(\frac{\pi}{2} - \psi_{t_2}^e(x) \right) \right) + \sum_{e \in \partial M \cap S} \left(\frac{\pi}{2} - \psi_t^e(x) \right) \\ &= \sum_{t \in S} \left(\pi + \frac{l^t(x)}{2} \right) + \sum_{e \in \partial S - \partial M} \left(\frac{\pi}{2} - \psi_t^e(x) \right) \\ &= \sum_{t \in S} \left(\pi - \frac{k^t(x)}{2} \right) + \sum_{e \in \partial S - \partial M} \left(\frac{\pi}{2} - \psi_t^e(x) \right). \end{aligned}$$

q.e.d

Now every edge in $\partial S - \partial M$ faces a zero on its S^c side in a triangle with $k^t(x) = 0$ (see figure 16 once again), so

$$\sum_{e \in \partial S - \partial M} \left(\frac{\pi}{2} - \psi_t^e(x) \right) = 0.$$

Similarly each triangle has zero curvature so from the above formula we have

$$\sum_{\{e \in S\}} \theta^e(x) = |S|\pi$$

violating condition (n_2) . So we have constructed a violation to (n_2) and $\Psi(x)$ cannot be in D as need.

q.e.d




Figure 17: Let $v \in V_\infty$ and call the vector in this figure f_v . Notice that this vector is $\sum_{e \in v} w_e$ (with the correct sign choices) hence a conformal transformation. For the purpose here, the most important observation is that constant multiples of this vector are precisely the conformal transformations involving v that preserve the condition that the angles at v remain zero. In general $\text{span}_{v \in V} \{f_v\}$ arises naturally as the set of transformations preserving each triangles curvature.

5 Generalizations and Comments

5.1 Corners, Cones and Cusps

Its worth noting that the proof of theorem 1 relies in no way on the restriction that $p_v(x) = 2\pi$ at internal vertices and $p^v(x) = \pi$ at boundary vertices. This assumption gives the simplest case, of producing hyperbolic surfaces with geodesic boundary. If this condition is relaxed to using any positive numbers, the same exact proof goes through to produce triangulations of surfaces with cornered boundary and cone singularities.

By far the most interesting case is when this condition is set to $p_v(x) = 0$ at some vertices, and we produce surfaces with cusps. For the discussion here let us suppose that this condition is placed at a set of interior vertices denoted V_∞ and that such vertices are isolated in the sense that no single triangle has two of its vertices in V_∞ . Notice $p_v(x) = 0$ forces us to make all the angles at v identically zero. Letting E_∞ and F_∞ denote the edges and faces with a vertex in V_∞ respectively, our space of angle systems, \mathbf{N} is restricted to $(0, \pi)^{3F - E_\infty}$. The conformal transformations are spanned by the w_e at edges in $E - E_\infty$ along with the vectors f_v from figure 17 for each $v \in V_\infty$.

On our new \mathbf{N} we can put on the same objective function H , though for faces in F_∞ the prism degenerates to the polyhedron in figure 18. The same exact same




Figure 18: Here is the notation for our degenerate prism, when $d = \{A_{i-1}, A_i, 0\}$ with A_k greater than zero. Notice the side of our original prism where the zero angle lives degenerates to zero length, and the pictured vertex v is on S^∞ .

arguments as in section 3 allow us to see that at a critical point the edges in $E - E_\infty$ fit together. The remaining edges are all infinite and so we can certainly glue the triangles of a flower at $v \in V_\infty$ together. To understand this situation, pick a cyclic order on the triangles this flower $\{t_1, \dots, t_n\}$ and denote the ordered pair of E_∞ edges of t_i as $\{e_{i-1}, e_i\}$. With this notation place the realizations of the $d^{t_i}(x)$ in the upper half plane as in figure 19.

At this point we need to know that at a critical point we satisfy the cusp producing condition as in figure 19; the trick to accomplishing this is to use the horocircle independence principle introduced in section 2.2 along with a "holonomy" argument. Notice that $l^{e_i}(u)$ is infinite. As usual by using a horoball cut off the length of $(l^{e_i}(u) - l^{e_{i-1}}(u))$ is well defined. This observation being true for each t_i independently allows us to use a simultaneous cut off of the whole realization in the upper-half plane as in figure 19.

Using the notation of figure 20 we see that e_0 and e_n matching up correctly is equivalent to

$$e_{n+1}^* - e_0^* = \sum_{i=1}^n (e_i^* - e_{i-1}^*) = 0.$$

(This observation is what I have referred to as an holonomy argument and it originally came up in this type of situation in Bragger's treatment of the Euclidean case, see [1]).

The usual angle formulas (from the proof of observation 1) hold for A_i^* and A_{i-1}^* , for exactly the same reasons. Just as in section 3.1 we will be computing dV via Schläfli's formula and will use any horosphere to cut off the vertex correspond-




Figure 19: In this figure we have placed the realizations of the triangles in our flower next to each other. When we glue the realization of e_0 to the realization of e_n we will typically not get a cusp, in fact such a gluing will usually produce a non-complete hyperbolic surface (see [6]). However, when the dotted horocircles in the figure agree, we may use an isometry represented by a Euclidean translation of the upper-half plane to identify e_0 and e_n and produce a cusp with piecewise geodesic boundary. The thick line is a horocircle which could be used to simultaneously cut all the infinite triangle sides.

ing to our $v \in V_\infty$ and the horospheres tangent to the specified H^2 to cut off the remaining "prism" vertices. Using the notation of figure 20 at a critical point we have

$$0 = dH(f_v) = \sum_{e \in v} (e_i^{**} - e_i^{**}).$$

However we can also see in figure 20 that $e^* = e^{**} + \ln(2)$, so this equation also implies the needed equation in the previous paragraph. Hence theorem 1 holds in the cusp case as well.

5.2 The Convex Case

Theorem 2 can of course be extended to the above mentioned cases as well. Much more importantly though, it can be extended to the entire convex case where $\theta^e(p) \in (0, \pi]$. This generalization is very useful to produce as corollaries other versions of the Thurston-Andreev theorem. The price is that there are new affine conditions placed on $p \in \mathbf{R}^E$ and the analogs to the linear arguments in section 4 become considerably more complicated. I'll state the theorem whose careful proof can be found in [3]. In order to articulate the new conditions we need certain snake and a loop concepts in a triangular decomposition.




Figure 20: The face of the degenerate prism containing the edge e_i is pictured, along with its horosphere cuts. For convenience e_i^* will denote both the cut off version of e_i 's realization with respect to the critical point's data and e^* 's length. Similarly for e_i^{**} , the degenerate prism's labeled edge. Notice by transforming this picture to the upper-half plane model we can immediately compute that $e^* = e^{**} + \ln(2)$.

Definition 9 A snake is a finite directed sequence of edges $\{e_i\}_{i=k}^l$ directed in the following sense: if $k < l$ we start with the edge e_k between t_{k-1} and t_k , then we require e_{k+1} to be one of the remaining edges on t_k . Then letting t_{k+1} be the other face associated to e_{k+1} we require e_{k+2} to be one of the other edges of t_{k+1} and so on until some tail edge e_l and tail face t_l are reached, and if $l < k$ we reverse the procedure and subtract from rather than add to the index. A loop is a snake $\{e_i\}_{i=l}^k$ where $e_k = e_l$ and $t_k = t_l$.

It is a condition on snakes and loops which allows one to articulate the remaining necessary conditions. It should be clear already how such objects are naturally born from the arguments in lemmas 7 and 8, in fact one can see the accordion in figure 15 and the vacuums in figure 16 for examples of a loop and snakes respectively. As defined snakes and loops can self intersect and there will be an infinite number of such objects, and it is worthwhile to first isolate a finite sub-set that does the job.

Definition 10 A set of edges $\{e_i\}_{i=k}^l$ is called embedded if $e_i \neq e_j$. A snake $\{e_i\}_k^l$ is said to double back on itself if we have a pair of non-empty sub-snakes with $\{e_i\}_m^n$ and $\{e_i\}_{k-m}^{k-n}$ containing the same edges. A barbell is a loop which doubles back on itself and such that $\{e_i\}_{i=k}^l / \{e_i\}_{i=m}^n$ is embedded. A balloon is a snake which doubles back on itself with $\{e_i\}_{i=k}^l / \{e_i\}_{i=m}^n$ embedded and such that

$$e_l = e_k.$$

With this terminology the remaining necessary conditions are

$$(n_3) \quad \sum_{i=k}^{l-1} \theta_i^e(p) < |k - l|\pi \quad \text{when } \{e_i\}_{i=k}^l \text{ is an embedded loop or barbell,}$$

and

$$(n_4) \quad \sum_{i=k}^l \theta_i^e(p) < (|k - l| + 1)\pi \quad \text{when } \{e_i\}_{i=k}^l \text{ is an embedded sake or balloon.}$$

With these conditions we can completely characterize the angles arising in convex ideal disk patterns.

Theorem 3 *That $p \in (0, \pi]^E$ and satisfies n_i for each i is necessary and sufficient for p to be equal to $\Psi(u)$ for some uniform angle system u . Furthermore this u is unique.*

5.3 Some Natural Questions

Here I list four cases where I think it would be particularly nice to attempt to apply these hyperbolic volume techniques.

1. **The Spherical Case** The use of hyperbolic volume in this paper to solve theorems 1 and 2 could conceivably be used to prove the analogous questions in the spherical case. The polyhedron to be used now becomes the twisted prism in figure 21 (also observed as the right object for this game by Peter Doyle). It is easy to see that the critical points of the hyperbolic volume objective function are once again precisely the uniform angle system; but the objective function fails to be convex. Can this objective function still be used to arrive at the analogous results? (See [3] for a more detailed account of this situation.)

2. **The Non-compact Case**

Analogous to theorems 2 and 1 also exist in non-compact situations. Let \mathbf{T} be a locally finite triangulation on a topological surface and let S be some set of triangles in \mathbf{T} . Given any finite $\hat{S} \subset S$ let $\partial_S \hat{S}$ be the subset set of \hat{S} boundary edges which are not S boundary edges. Denote as (\hat{n}_2) the condition that for any set S there is some $\hat{S} \subset S$ such that

$$\sum_{e \in \hat{S}} \theta^e(p) - \pi|\hat{S}| > \pi|\partial_S \hat{S}|.$$




Figure 21: Fix a point p in H^3 viewed in the Poincare Model. Imagine S^∞ is given the unit sphere's induced metric from \mathbf{R}^3 . Given a geodesic triangle on with a hyperbolic geodesic to p . The polyhedron of interest is the convex hull of these geodesic segments.

Note this condition is equivalent to (n_2) when \mathbf{T} is finite.

Let condition (\hat{n}_4) be that for any **infinite** snake $\{e_i\}$ there exist N_i such that

$$\sum_{i=N_1}^{N_2} \psi^{e_i}(p) > \pi.$$

Note this condition is automatically satisfied in the finite case as well.

Let a geodesic triangulation satisfy having all its angles in $(\epsilon, \pi]$ and all its k^t in $[-\pi, -\epsilon)$ for some $\epsilon > 0$. Then it is straightforward to verify that (\hat{n}_2) and (\hat{n}_4) are in fact necessary. The following corollary to theorem 2 guarantees these condition are always sufficient.

Corollary 3 *Given $p \in (0, \pi)^{(E-\partial E)} \times (0, \frac{\pi}{2})^{\partial E}$ satisfying (n_1) , (\hat{n}_2) , and (\hat{n}_4) , then p is an ideal disk pattern.*

Two question immediately arise. First would be nice to know conditions guaranteeing when there is a solution forming a complete surface. Secondly it would be nice to to understand the uniqueness, or more likely the moduli of space of possible solutions.

3. **The Topological Case** Given a topological triangular decomposition \mathbf{T} theorem 2 tells whether a given set of intersection angle data can be geometrically realized as an Delaunay ideal disk pattern. It would be nice to answer:




Figure 22: A topological triangulation guaranteed to have a geometric realization as a Delaunay triangulation. This guarantee still hold when we glue together copies of the pictured pairs of pants.

for which \mathbf{T} does consistent intersection angle data exist. There are some several easy answers. For example letting $\{n_1, n_2, n_3\}$ be the degrees of the vertices of any triangle, if we have

$$\frac{1}{n_1} + \frac{1}{n_2} + \frac{1}{n_3} < \frac{1}{2}$$

and

$$0 < \frac{1}{n_1} + \frac{1}{n_2} - \frac{1}{n_3},$$

then the angles $\frac{2\pi}{n_i}$ satisfies theorem 1 (see figure 22). It would be nice to find an identification procedure that worked for any triangulation.

Furthermore as observed in section 3.2 H is convex through out \mathbf{N} . A critical point here is the maximal volume associated polyhedron among all possible realizations. Such a critical point turns out to be a particularly symmetric triangulation, where triangles t_1 and t_2 sharing an edge e with edges given by $\{a_i, b_i, e\}$ will satisfies $a_1^* + b_1^* = a_2^* + b_2^*$. In the presence of such a critical point we then have a canonical realization of \mathbf{T} , and it would be nice know for which \mathbf{T} does this realization exists?

4. **The Non-ideal Case** It would be nice to extend this use of hyperbolic volume to the non-ideal cases, and in particular prove theorems 1 and 2 in these cases.

In the sub-ideal case (when the vertices of the hyperbolic polyhedra associated to the ideal disk pattern are finite) the natural hyperbolic objective function created using the corresponding finite prisms can once again by Schläfli's formula be seen to have uniform critical points in a "conformal class". However to get started one must first construct the perpendicular edge lengths corresponding to each vertex (the $(ab)^*$, $(ac)^*$, and $(bc)^*$ in figure 8), so that the angles sum to 2π at each vertex. In the process the objective function becomes de-localized, making boundary control and concavity difficult to verify, and the corresponding theorem 1 difficult to get a handle on.

It is worth noting one can use theorem 3 to achieve disk packing and super-ideal versions of theorems 3, but one needs to retriangulate a bit and the process feels a bit synthetic. It would be nice to have direct hyperbolic volume techniques in this case as well.

References

- [1] Walter Brägger. Kreispackungen und triangulation. *Ens. Math*, 38:201–217, 1992.
- [2] Marvin Jay Greenberg. *Euclidean and Non-euclidean Geometries*. W. H. Freeman and Co., San Francisco, second edition, 1972.
- [3] G. Leibon. *Random Delaunay Triangulations, the Thurston-Andreev Theorem and Metric Uniformization*. PhD thesis, UCSD, 1999.
- [4] I. Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Ann. of Math (2)*, 139:553–580, 1994.
- [5] I. Rivin. A characterization of ideal polyhedra in hyperbolic 3-space. *Ann. of Math.(2)*, 143:51–70, 1996.
- [6] William P. Thurston. *Three-Dimensional Geometry and Topology*. The Geometry Center, University of Minnesota, draft edition, 1991.